

Introduction to Apache Kafka for Developers

- **Course: Introduction to Apache Kafka for Developers / Working with Apache Kafka (TTDS6760)**
 - **Duration:** 2 days, or four half days
 - **Skill Level:** This introductory-level course is geared for skilled, experienced Java Developers and architects with Java development background who are **new to Kafka**. This course is NOT for non-developers. Students without development background can follow along with or cut and paste the labs.
 - **Hands-on:** This hands-on course combines engaging instructor-led presentations and practical demonstrations with extensive programming exercises, challenge labs, use case exploration and engaging group activities. Student machines are required.
 - **Format:** This course can be delivered for your team or organization **online live (virtual)**, **onsite in-person** or across our robust **blended learning platform (LXP)**.
 - **Public Schedule:** This course is currently available on our Public Open Enrollment Schedule.
 - **Customizable:** This course agenda, topics, labs, hours and delivery modalities can be adjusted to target your specific training skills objectives, tools and learning goals. Please ask for details.
-

Overview

Apache Kafka is a real-time data pipeline processor. Its high-scalability, fault tolerance, execution speed, and fluid integrations are some of the key hallmarks that make it an integral part of many Enterprise Data architectures.

Discover the future of data streaming with our hands-on, 2-day Apache Kafka training course designed specifically for experienced Java developers. Become an expert in harnessing the potential of fast data and streaming systems and learn how to navigate the complexities of modern streaming architectures. With practical labs and real-world examples, you'll be immersed in the cutting-edge world of Kafka and Java development.

Throughout the course you'll explore the ins and outs of Apache Kafka and learn how it compares to other queue systems like JMS and MQ. You'll learn about Kafka's unique architecture and understand how to effectively produce and consume messages with Kafka & Zookeeper. Through hands-on labs, you'll gain experience in scaling Kafka, navigating multiple data centers, and implementing disaster recovery solutions, while exploring essential Kafka utilities.

You'll also learn the powerful Kafka APIs and become proficient in configuration parameters, Producer and Consumer APIs, as well as advanced features such as message compression and offset management. Gain hands-on with Kafka, including benchmarking Producer send modes, comparing compression schemes, and managing offsets. Experience real-world applications like Clickstream processing to solidify your expertise. Then you'll round off your Kafka journey with an in-depth look at the Kafka Streams API, monitoring, and troubleshooting techniques. Learn how to optimize your Kafka deployment with best practices for hardware selection, cluster sizing, and Zookeeper settings.

By the end of this course you'll be equipped with the core skills required to tackle your next Kafka project with confidence.

NOTE: Experienced developers who want a deeper hands-on dive into Kafka might consider the **Exploring Kafka (TTDS6764)** four day course as an alternative.

Learning Objectives

Working in a hands-on learning environment you'll learn to:

- Implement and configure Apache Kafka effectively, demonstrating a deep understanding of its unique architecture, core concepts, and the differences between Kafka and other queue systems (JMS/MQ).

- Utilize Kafka APIs proficiently, including the Producer and Consumer APIs, and apply advanced techniques such as message compression, offset management, and Producer send modes.
- Design and develop streaming applications using the Kafka Streams API, performing complex operations like transformations, filters, joins, and aggregations, while working with KStream, KTable, and KStore concepts.
- Monitor and troubleshoot Kafka deployments, identifying performance bottlenecks, addressing common issues, and employing best practices for hardware selection, cluster sizing, partition sizing, and Zookeeper settings.
- Apply the skills and knowledge acquired throughout the course to real-world scenarios, showcasing the ability to develop, deploy, and optimize Kafka-based streaming applications for a variety of use cases.

Audience: Who Should Attend?

This course is geared for experienced Java Developers and architects with Java development background who are new to Kafka. **This course is not for non-developers.**

In order to be successful in this course, and to participate in the hands-on labs, you should possess:

- Basic Java programming skills; practical Java development background.
- Reasonable experience working with databases
- Basic Linux skills and the ability to work from the Linux command line
- Basic knowledge of Linux editors (such as VI / nano) for editing code.

Take Before: Students should have practical skills equivalent to or should have attended the following course(s) as a pre-requisite:

- TT2104: Fast Track to Core Java Programming for OO Developers (C+, C#, etc.)

Next Steps / Follow-on Courses: We offer a wide variety of follow-on courses for next-level Kafka, Apache tools, data science, machine learning, DevOps, Java development skills and more. Please see our **Learning Journey & Career Experience Paths** for options based on your specific role and goals.

Course Topics / Agenda

We can collaborate with you to tune this course and level of coverage to target the skills you need most. Course agenda, topics and labs are subject to adjust during live delivery in response to student skill level, interests and participation.

Introduction to Streaming Systems

- Understanding Fast data
- Streaming terminologies
- Understanding at-least-once / at-most-once / exactly-once processing patterns
- Popular streaming architectures
- Lambda architecture
- Streaming platforms overview
- Lab: Hands-on first look at Kafka

- Kafka & Zookeeper
- Producing messages
- Consuming messages
- Consumers, Consumer Groups
- Message retention
- Scaling Kafka
- Kafka across multiple data centers and disaster recovery
- Lab: Getting Kafka up and running
- Lab: Using Kafka utilities

- Producer send modes
- Message compression
- Commits , Offsets, Seeking
- Managing offsets - auto commit / manual commit
- Lab: Writing Producer / Consumer
- Lab: Benchmarking Producer send modes
- Lab: Comparing compression schemes
- Lab: Managing offsets
- Lab: Clickstream processing

Introducing Kafka

- Comparing Kafka with other queue systems (JMS / MQ)
- Kafka Architecture
- Kaka concepts: Messages, Topics, Partitions, Brokers, Producers, commit logs

Using Kafka APIs

- Configuration parameters
- Producer API - sending messages to Kafka
- Consumer API - consuming messages from Kafka

Kafka Streams API

- Introduction to Kafka Streams library
- Features and design

- Streams concepts: KStream / KTable / KStore
- Streaming operations (transformations, filters, joins, aggregations)
- Using Streams API: foreach / filter / map / groupby
- Lab: Kafka Streaming APIs

Monitoring & Troubleshooting Kafka

- Monitoring tools overview
- Monitoring Kafka
- Cluster level and host level monitoring
- Identifying performance bottlenecks
- Troubleshooting common Kafka issues

Bonus Content / Time Permitting

Kafka Best Practices

- Avoiding common mistakes
- Hardware selection
- Cluster sizing
- Partition sizing
- Zookeeper settings
- Compression and batching
- Message sizing
- Monitoring and instrumenting
- Troubleshooting

Setup Made Simple with our Robust Learning Experience Platform (LXP)

All course software (limited versions, for course use only), knowledge checks, digital courseware files or course notes, labs / data sets and solutions, live coaching support channels (as applicable) and rich extended learning / post training resources are provided for you in our “easy access / no install required” high-speed **Learning Experience Platform (LXP)** remote lab and content environment. Our tech team works with every student to ensure everyone is set up with working access and ready to go prior to every course start date, ensuring a smooth delivery and great hands-on experience.

For More Information

For more information about our dedicated skills-focused training services (instructor-led, self-paced or blended), collaborative coaching services, robust Learning Experience Platform (LXP) solutions, Career Experiences, public course schedule, partner programs, courseware licensing options or to see our complete list of course offerings, training solutions and special offers please visit us at www.triveratech.com, email Info@triveratech.com or call us toll free at **844-475-4559**. Our pricing and services are always satisfaction guaranteed.

TRIVERA TECHNOLOGIES • Collaborative IT Training, Coaching & Skills Development Solutions
www.triveratech.com • toll free +1-844-475-4559 • Info@triveratech.com • Twitter TriveraTech

ONSITE, ONLINE & BLENDED TRAINING SOLUTIONS • PUBLIC / OPEN ENROLLMENT COURSES
 LEARNING EXPERIENCE PLATFORM (LXP) • COACHING / MENTORING • ASSESSMENTS • CONTENT LICENSING & DEVELOPMENT
 LEARNING PLAN DEVELOPMENT • SKILLS IMMERSION PROGRAMS / RESKILLING / NEW HIRE / BOOT CAMPS
 PARTNER & RESELLER PROGRAMS • CORPORATE TRAINING MANAGEMENT • VENDOR MANAGEMENT SERVICES

Trivera Technologies is a Woman-Owned Small-Business Firm

